

WHAT IS CLAIMED IS:

1. A computer network file system, comprising:

first metadata managed primarily by a first file server operably connected to a network fabric, said first metadata comprising first file location information, said first file location information comprising at least one server id; and

second metadata managed primarily by a second file server operably connected to the network fabric, said second metadata comprising second file location information, said second file location information comprising at least one server identifier, said first metadata and said second metadata configured to allow a requestor to locate files stored by said first file server and files stored by said second file server in a directory structure that spans said first file server and said second file server.

2. The computer network file system of Claim 1, said first file location information comprising a disk identifier.

3. The computer network file system of Claim 1, wherein said first metadata comprises first directory information that describes at least a portion of said directory structure relating to directories stored on said first file server and second metadata that describes at least a portion of said directory structure relating to directories stored on said second file server, said first directory information comprising location information for said second directory information, said location information comprising a server identifier that identifies said second file server.

4. The computer network file system of Claim 3, wherein said first metadata comprises a root directory.

5. The computer network file system of Claim 3, wherein said location information further comprises a disk identifier that identifies a disk drive attached to a file server corresponding to said server identifier.

6. The computer network file system of Claim 2, wherein said first metadata comprises at least one Gnid-string.

7. The computer network file system of Claim 6, wherein a one-to-one correspondence exists between said at least one Gnid-string and a directory owned by said first file server.

8. The computer network file system of Claim 6, wherein said Gnid-string comprises a collection of gnids.

9. The computer network file system of Claim 8, wherein each of said gnids comprises information for locating a specified gnode.

5 10. The computer network file system of Claim 9, wherein said information for locating a specified gnode comprises a pointer to said specified gnode.

11. The computer network file system of Claim 9, wherein said specified gnode comprises file attributes for a file corresponding to said gnode.

10 12. The computer network file system of Claim 11, wherein said file attributes include at least one of a file id, a file access time, a file creation time, and a file modification time.

13. The computer network file system of Claim 9, wherein said specified gnode comprises information for locating a first gee of a plurality of gees corresponding to said gnode.

15 14. The computer network file system of Claim 13, wherein said plurality of gees comprises gnode gees and data gees.

15. The computer network file system of Claim 14, wherein each of said gnode gees comprises information to specify an extent.

20 16. The computer network file system of Claim 14, wherein each of said data gees comprises information to specify a first logical disk block and information to specify a disk that contains said first logical block.

17. The computer network file system of Claim 13, wherein said plurality of gees further comprises parity gees.

25 18. The computer network file system of Claim 17, wherein each of said parity gees comprises information regarding location of parity data for one or more preceding data gees in said plurality of gees.

19. The computer network file system of Claim 13, wherein a parity group comprises a first set of one or more data gees and an associated parity gee.

30 20. The computer network file system of Claim 19, wherein each data gee identifies a block of data and said parity gee identifies a parity block.

21. The computer network file system of Claim 20, wherein each block of data and parity in said parity group is stored on a separate disk drive such that no single disk drive contains data from two blocks said parity group.

22. The computer network file system of Claim 19, wherein a size of a first parity group is independent of a size of a second parity group.

23. The computer network file system of Claim 1, further comprising an intent log maintained by said first file server.

24. The computer network file system of Claim 1, wherein said network fabric comprises a Fibre channel network.

25. The computer network file system of Claim 1, wherein said network fabric comprises an ethernet network.

26. The computer network file system of Claim 1, wherein said network fabric comprises an asynchronous transfer mode network.

27. The computer network file system of Claim 1, wherein said network fabric comprises a first Fibre channel network and wherein said first file server communicates with one or more disk drives using a second Fibre channel network.

28. The computer network file system of Claim 1, wherein files stored by said first file server and files stored by said second file server are located in a hierarchical directory structure having a common root directory, said first metadata comprising directory information that describes directories stored on said first file server, said second metadata comprising directory information that describes directories stored on said second file server, said first directory information comprising location information for finding said second directory information, said location information comprising a server identifier that identifies said second file server.

29. The computer network file system of Claim 1, wherein files stored by said first file server and files stored by said second file server are located in a hierarchical directory structure having a common root directory, said first file metadata comprising metadata for locating files stored on said second file server.

30. The computer network file system of Claim 1, wherein first metadata describes directories stored on said first file server, and where said second metadata describes directories stored on said second file server, said first metadata comprising

location information for locating said second metadata, said location information comprising a server identifier.

31. The computer network file system of Claim 30, wherein said first metadata comprises a root directory.

32. The computer network file system of Claim 30, wherein at least said first metadata comprises file attributes for one or more files stored by said first file server.

33. The computer network file system of Claim 30, wherein at least said first metadata comprises information that specifies logical disk blocks of a selected file and information to specify which disks contain said logical disk blocks.

34. The computer network file system of Claim 30, wherein said metadata identifies data blocks and parity blocks corresponding to said data blocks.

35. The computer network file system of Claim 30, wherein said metadata identifies parity groups, said parity groups comprising a plurality of information blocks, said information blocks comprising one or more data blocks, said information blocks further comprising a parity block, each of said information blocks stored on a different disk drive.

36. The computer network file system of Claim 1, wherein a size of a first distributed parity group of a file is independent of a size of a second distributed parity group of said file.

37. A method for storing data in a computer network, comprising:
creating first file system metadata on a first file server operably connected to a network fabric, said first file system metadata describing at least files and directories stored by said first file server;

creating second file system metadata on a second file server connected to said network fabric, said second file system metadata describing at least files and directories stored by said second file server, said first file system metadata and said second file system metadata comprising directory information that spans said first file server and said second file server, said directory information configured to allow a requestor to find a location of a first file catalogued in said directory information without prior knowledge as to a server location of said first file.

38. The method of Claim 37, wherein said first file system metadata comprises a root directory.

39. The method of Claim 37, wherein said first file system metadata comprises at least one Gnid-string.

5 40. The method of Claim 39, wherein a correspondence exists between said at least one Gnid-string and a directory corresponding to said first file server.

41. The method of Claim 39, wherein said Gnid-string comprises a collection of gnids.

10 42. The method of Claim 41, wherein each of said gnids comprises information for locating a specified gnode.

43. The method of Claim 42, wherein said information for locating a specified gnode comprises a pointer to said specified gnode.

44. The method of Claim 42, wherein said specified gnode comprises file attributes for a file corresponding to said gnode.

15 45. The method of Claim 44, wherein said file attributes include at least one of a file id, a file access time, a file creation time, and a file modification time.

46. The method of Claim 42, wherein said specified gnode comprises information for locating a first gee of a plurality of gees corresponding to said gnode.

20 47. The method of Claim 46, wherein said plurality of gees comprises gnode gees and data gees.

48. The method of Claim 47, wherein each of said gnode gees comprises information to specify a logical block extent.

25 49. The method of Claim 47, wherein each of said data gees comprises information to specify a first logical disk block and information to specify a disk that contains said first logical block.

50. The method of Claim 49, further comprising defining a plurality of parity gees.

30 51. The method of Claim 50, wherein each of said parity gees comprises information regarding location of parity data for one or more preceding data gees in said plurality of gees.

52. The method of Claim 37, further comprising defining at least one parity group having a first parity group size, said at least one parity group comprising a parity block and one or more data blocks.

53. The method of Claim 52, further comprising storing each of said data blocks and said parity block on different disk drives.

54. The method of Claim 53, further comprising defining at least one parity group having a second parity group size.

55. The method of Claim 37, further comprising mirroring said first file system directory information on said second file server.

56. The method of Claim 37, wherein said file system metadata describes at least a portion of a hierarchical directory structure, said hierarchical directory structure spanning said first file server and said second file server, said first file system metadata comprising directory information that describes directories stored on said first file server, and wherein said second file system metadata describes at least a portion of said hierarchical directory structure comprising directories stored on said second file server, said first directory information comprising location information for finding said second directory information, said location information identifying said second file.

57. The method of Claim 37, wherein said first metadata describes directories stored on said first file server, said first metadata comprising location information for locating said second metadata, said location information comprising a server identifier, said first metadata and said second metadata describing at least a portion of a hierarchical directory structure that spans both said first file server and said second file server.

58. The method of Claim 57, wherein said first metadata comprises a root directory.

59. The method of Claim 57, wherein at least said first metadata comprises file attributes for one or more files stored by said first file server.

60. The method of Claim 57, wherein at least said first metadata comprises information to specify a selected logical disk block of a selected file and information to specify a disk that contains said selected logical block.

61. The method of Claim 57, wherein said metadata identifies data blocks and parity blocks corresponding to said data blocks.

62. The method of Claim 37, further comprising performing a directory lookup by sending a directory lookup request to said first file server, said first file server returning a directory file handle that points to a directory on said second file server.

63. A computer network file system, comprising:

a first file server operably connected to a network fabric;

means for creating files stored on said first file server or a second file server; and

means for locating files stored by said first file server and files stored by said second file server by traversing a directory structure that spans at least a said first file server and said second file server.

64. The computer network file system of Claim 63, further comprising means for describing parity groups.